What is claimed is:

1. A method of adaptive direct volume rendering, comprising

fragmenting a sampled 3-D dataset of a scalar field into a plurality of sub-volumes of different sizes, each sub-volume associated with a set of data value parameters characterizing the data value distribution of the scalar field within the sub-volume;

defining an opacity transfer function that is dependent upon data values of the scalar field and an illumination model;

selectively casting a plurality of rays from a 2-D image plane towards the sampled dataset, each ray having an initial ray energy and a cross-section;

for each ray cast from a selected location on the 2-D image plane,

selecting a subset of the plurality of sub-volumes for interacting with the ray;

estimating the ray energy reflected by each sub-volume of the subset using the opacity transfer function and the illumination model; and

summing the reflected ray energy as a pixel value at the selected location on the 2-D image plane; and

estimating pixel values at other locations on the 2-D image plane using the pixel values at the selected locations.

2. The method of claim 1, wherein

the step of fragmenting a sampled 3-D dataset includes

fragmenting the 3-D dataset into eight sub-volumes; and

for each sub-volume, recursively fragmenting it into eight smaller sub-volumes until the size of the smallest sub-volumes reaches a predefined size limit.

3. The method of claim 2, wherein the predefined size limit is a sub-volume comprising 2x2x2 3-D cells and the eight corners of each cell are associated with eight data values of the scalar field.

4. The method of claim 3, wherein the data value at a location within the cell is tri-linearly interpolated using the eight data values at the eight corners of the cell.

5. The method of claim 1, wherein the set of parameters include a maximum, an average value, and a minimum data value of the scalar field within the sub-volume.

6. The method of claim 1, further comprises

constructing an octree comprising a root node, a plurality of intermediate nodes, and a plurality of leaf nodes;

associating the root node with the 3-D dataset;

associating each of the plurality of leaf nodes with a smallest sub-volume from the plurality of sub-volumes; and

associating each of the plurality of intermediate nodes with a sub-volume from the plurality of sub-volumes that is larger than the smallest sub-volume.

7. The method of claim 1, wherein

the step of casting a plurality of rays from a 2-D image plane includes

subdividing the 2-D image plane into a plurality of sub-planes;

for each of the plurality of sub-planes,

casting four rays from the four corners of the sub-plane and estimating a pixel value at each corner;

calculating a maximum pixel value variation within the sub-plane; and

recursively subdividing the sub-plane into multiple child sub-planes of smaller sizes by casting a ray from the center of the sub-plane until the maximum pixel value variation of the sub-plane is below a predefined imaging error threshold.

8. The method of claim 7, wherein the maximum pixel value variation within the sub-plane is defined as the maximum deviation of pixel values at the four corners of the sub-plane from an average pixel value of the sub-plane.

9. The method of claim 7, wherein the predefined imaging error threshold is modulated by an image rendering speed provided by a user, a distance to an edge of an object embedded in the 3-D dataset, and a difference between a pixel value estimated from an adaptive ray casting and a pixel value estimated from a bi-linear interpolation.

10. The method of claim 1, wherein

the step of selecting a subset of the plurality of sub-volumes for interacting with the ray includes

identifying a largest sub-volume along the ray path and its corresponding maximum and minimum data values;

checking if the opacity transfer function varies monotonically between the maximum and minimum scalar field values;

if the function does not vary monotonically, recursively

identifying a smaller sub-volume along the ray path and its corresponding maximum and minimum data values; and

checking if the opacity transfer function varies monotonically between the maximum and minimum scalar field values of the smaller sub-volume; and

if the function does vary monotonically, calculating the amount of ray energy reflected by the sub-volume during its interaction with the ray.

11. The method of claim 10, wherein two lookup tables are constructed for the opacity transfer function such that a forward lookup table contains the data value difference to a nearest local extreme of the opacity transfer function along the data value increasing direction and a backward lookup table contains the data value difference to a nearest local extreme of the opacity transfer function along the data value decreasing direction.

12. The method of claim 11, wherein if the maximum data value of the sub-volume is smaller than the summation of the minimum data value of the sub-volume and its corresponding data value difference stored in the forward lookup table or the minimum data value of the sub-volume is larger than the difference between the maximum data value of the sub-volume and its corresponding data value difference stored in the backward lookup table, the opacity transfer function varies monotonically between the minimum and maximum data values.

13. The method of claim 1, wherein

the step of estimating the ray energy reflected by each sub-volume of the subset includes

estimating a maximum energy differential of the sub-volume;

comparing the maximum energy differential against a predefined energy error threshold;

if the maximum energy differential is above the predefined energy error threshold, recursively

selecting a smaller sub-volume along the ray path; and

estimating a new maximum energy differential of the smaller sub-volume; and

if the maximum energy differential is below the predefined energy error threshold, calculating the amount of ray energy reflected by the sub-volume using the illumination model.

14. The method of claim 13, wherein the maximum energy differential depends on the opacity transfer function and the maximum, average, and minimum data values of the sub-volume.

15. The method of claim 13, wherein the amount of ray energy reflected by the sub-volume depends on the length of ray path within the sub-volume, the opacity transfer function within the sub-volume, the average scalar field value of the sub-volume, and the local gradient vector of scalar field within the sub-volume.

16. The method of claim 13, wherein if the sub-volume is a smallest sub-volume comprising 2x2x2 3-D cells, the smaller sub-volume is a 3-D cell within the smallest sub-volume, and if the opacity function does not vary monotonically within the cell, the 3-D cell is further divided into multiple sub-cells until the dimension of a smallest sub-cell reaches the cross-section of the ray.

17. The method of claim 13, wherein if the sub-volume is a smallest sub-volume comprising 2x2x2 3-D cells, the smaller sub-volume is a 3-D cell within the smallest sub-volume, and if the opacity function varies monotonically within the cell, the maximum energy differential of the 3-D cell is estimated by dividing the maximum energy differential of the sub-volume by 2.

18. The method of claim 13, wherein if the sub-volume is a smallest sub-volume comprising 2x2x2 3-D cells, the smaller sub-volume is a 3-D cell within the smallest sub-volume, and if the opacity transfer function varies monotonically within the cell and an iso-surface exists in the 3-D cell, the maximum energy differential of the 3-D cell is calculated using the eight data values at the corners of the 3-D cell and the opacity transfer function.

19. The method of claim 13, wherein the predefined energy error threshold is modulated by an image rendering speed specified by a user and a zoom factor in the case of parallel projection or a perspective angle and a perspective distance between the image plane and the 3-D dataset in the case of perspective projection.

20. The method of claim 1, wherein

the step of estimating pixel values at other locations on the 2-D image plane includes for each location,

selecting four pixel values associated with four ray origins surrounding the location; and

bi-linearly interpolating a pixel value at the location using the four pixel values.

21. An adaptive direct volume rendering system, comprising:

one or more central processing units for executing programs;

a user interface for receiving a plurality of volume rendering parameters; and

an adaptive volume rendering engine module executable by the one or more central processing units, the module comprising:

instructions for fragmenting a sampled 3-D dataset of a scalar field into a plurality of sub-volumes of different sizes, each sub-volume associated with a set of data value parameters characterizing the data value distribution of the scalar field within the sub-volume;

instructions for defining an opacity transfer function that is dependent upon data values of the scalar field and an illumination model;

instructions for selectively casting a plurality of rays from a 2-D image plane towards the sampled dataset, each ray having an initial ray energy and a cross-section;

for each ray launched from a selected location on the 2-D image plane,

instructions for selecting a subset of the plurality of sub-volumes for interacting with the ray;

instructions for estimating the ray energy reflected by each sub-volume of the subset using the opacity transfer function and the illumination model; and

instructions for summing the reflected ray energy as a pixel value at the selected location on the 2-D image plane; and

instructions for estimating pixel values at other locations on the 2-D image plane using the pixel values at the selected locations.

22. The system of claim 21, wherein

the step of fragmenting a sampled 3-D dataset includes

fragmenting the 3-D dataset into eight sub-volumes; and

for each sub-volume, recursively fragmenting it into eight smaller sub-volumes until the size of the smallest sub-volumes reaches a predefined size limit.

23. The system of claim 22, wherein the predefined size limit is a sub-volume comprising 2x2x2 3-D cells and the eight corners of each cell are associated with eight data values of the scalar field.

24. The system of claim 23, wherein the data value at any location within the cell is tri-linearly interpolated using the eight data values at the eight corners of the cell.

25. The system of claim 21, wherein the set of parameters include a maximum, an average value, and a minimum data value of the scalar field within the sub-volume.

26. The system of claim 21, further comprises

instructions for constructing an octree comprising a root node, a plurality of intermediate nodes, and a plurality of leaf nodes;

instructions for associating the root node with the 3-D dataset;

instructions for associating each of the plurality of leaf nodes with a smallest sub-volume from the plurality of sub-volumes; and

instructions for associating each of the plurality of intermediate nodes with a sub-volume from the plurality of sub-volumes that is larger than the smallest sub-volume.

27. The system of claim 21, wherein

the instructions for selectively casting a plurality of rays from a 2-D image plane include

subdividing the 2-D image plane into a plurality of sub-planes;

for each of the plurality of sub-planes,

casting four rays from the four corners of the sub-plane and estimating a pixel value at each corner;

calculating a maximum pixel value variation within the sub-plane; and

recursively subdividing the sub-plane into multiple child sub-planes of smaller sizes by casting a ray from the center of the sub-plane until the maximum pixel value variation of the sub-plane is below a predefined imaging error threshold.

28. The system of claim 27, wherein the maximum pixel value variation within the sub-plane is defined as the maximum deviation of pixel values at the four corners of the sub-plane from an average pixel value of the sub-plane.

29. The system of claim 27, wherein the predefined imaging error threshold is modulated by an image rendering speed provided by a user, a distance to an edge of an object embedded in the 3-D dataset, and a difference between a pixel value estimated from an adaptive ray casting and a pixel value estimated from a bi-linear interpolation.

30. The system of claim 21, wherein

the instructions for selecting a subset of the plurality of sub-volumes for interacting with the ray include

identifying a largest sub-volume along the ray path and its corresponding maximum and minimum data values;

checking if the opacity transfer function varies monotonically between the maximum and minimum scalar field values;

if the function does not vary monotonically, recursively

identifying a smaller sub-volume along the ray path and its corresponding maximum and minimum data values; and

checking if the opacity transfer function varies monotonically between the maximum and minimum scalar field values of the smaller sub-volume; and

if the function does vary monotonically, calculating the amount of ray energy reflected by the sub-volume during its interaction with the ray.

31. The system of claim 30, wherein two lookup tables are constructed for the opacity transfer function such that a forward lookup table contains the data value difference to a nearest local extreme of the opacity transfer function along the data value increasing direction and a backward lookup table contains the data value difference to a nearest local extreme of the opacity transfer function along the data value decreasing direction.

32. The system of claim 31, wherein if the maximum data value of the sub-volume is smaller than the summation of the minimum data value of the sub-volume and its corresponding data value difference stored in the forward lookup table or the minimum data value of the sub-volume is larger than the difference between the maximum data value of the sub-volume and its corresponding data value difference stored in the backward lookup table,

the opacity transfer function varies monotonically between the minimum and maximum data values.

33. The system of claim 21, wherein

the instructions for estimating the ray energy reflected by each sub-volume of the subset include

estimating a maximum energy differential of the sub-volume;

comparing the maximum energy differential against a predefined energy error threshold;

if the maximum energy differential is above the predefined energy error threshold, recursively

selecting a smaller sub-volume along the ray path; and

estimating a new maximum energy differential of the smaller sub-volume; and

if the maximum energy differential is below the predefined energy error threshold, calculating the amount of ray energy reflected by the sub-volume using the illumination model.

34. The system of claim 33, wherein the maximum energy differential depends on the opacity transfer function and the maximum, average, and minimum data values of the sub-volume.

35. The system of claim 33, wherein the amount of ray energy reflected by the sub-volume depends on the length of ray path within the sub-volume, the opacity transfer function within the sub-volume, the average scalar field value of the sub-volume, and the local gradient vector of scalar field within the sub-volume.

36. The system of claim 33, wherein if the sub-volume is a smallest sub-volume comprising 2x2x2 3-D cells, the smaller sub-volume is a 3-D cell within the smallest sub-volume, and if the opacity function does not vary monotonically within the cell, the 3-D cell is further divided into multiple sub-cells until the dimension of a smallest sub-cell reaches the cross-section of the ray.

37. The system of claim 33, wherein if the sub-volume is a smallest sub-volume comprising 2x2x2 3-D cells, the smaller sub-volume is a 3-D cell within the smallest sub-volume, and if the opacity function varies monotonically within the cell, the maximum

energy differential of the 3-D cell is estimated by dividing the maximum energy differential of the sub-volume by 2.

38. The system of claim 33, wherein if the sub-volume is a smallest sub-volume comprising 2x2x2 3-D cells, the smaller sub-volume is a 3-D cell within the smallest sub-volume, and if the opacity transfer function varies monotonically within the cell and an iso-surface exists in the 3-D cell, the maximum energy differential of the 3-D cell is calculated using the eight data values at the corners of the 3-D cell and the opacity transfer function.

39. The system of claim 33, wherein the predefined energy error threshold is modulated by an image rendering speed specified by a user, and a zoom factor in the case of parallel projection or a perspective angle and a perspective distance between the image plane and the 3-D dataset in the case of perspective projection.

40. The system of claim 21, wherein

the instructions for estimating pixel values at other locations on the 2-D image plane include

for each location,

selecting four pixel values associated with four ray origins surrounding the location; and

bi-linearly interpolating a pixel value at the location using the four pixel values.

41. A computer program product for use in conjunction with a computer system, the computer program product comprising a computer readable storage medium and a computer program mechanism embedded therein, the computer program mechanism comprising:

instructions for fragmenting a sampled 3-D dataset of a scalar field into a plurality of sub-volumes of different sizes, each sub-volume associated with a set of data value parameters characterizing the data value distribution of the scalar field within the sub-volume;

instructions for defining an opacity transfer function that is dependent upon data values of the scalar field and an illumination model;

instructions for selectively casting a plurality of rays from a 2-D image plane towards the sampled dataset, each ray having an initial ray energy and a cross-section;

for each ray launched from a selected location on the 2-D image plane,

instructions for selecting a subset of the plurality of sub-volumes for interacting with the ray;

instructions for estimating the ray energy reflected by each sub-volume of the subset using the opacity transfer function and the illumination model; and

instructions for summing the reflected ray energy as a pixel value at the selected location on the 2-D image plane; and

instructions for estimating pixel values at other locations on the 2-D image plane using the pixel values at the selected locations.

42. The computer program product of claim 41, wherein

the step of fragmenting a sampled 3-D dataset includes

fragmenting the 3-D dataset into eight sub-volumes; and

for each sub-volume, recursively fragmenting it into eight smaller sub-volumes until the size of the smallest sub-volumes reaches a predefined size limit.

43. The computer program product of claim 42, wherein the predefined size limit is a sub-volume comprising 2x2x2 3-D cells and the eight corners of each cell are associated with eight data values of the scalar field.

44. The computer program product of claim 43, wherein the data value at any location within the cell is tri-linearly interpolated using the eight data values at the eight corners of the cell.

45. The computer program product of claim 41, wherein the set of parameters include a maximum, an average value, and a minimum data value of the scalar field within the sub-volume.

46. The computer program product of claim 41, further comprises

instructions for constructing an octree comprising a root node, a plurality of intermediate nodes, and a plurality of leaf nodes;

instructions for associating the root node with the 3-D dataset;

instructions for associating each of the plurality of leaf nodes with a smallest sub-volume from the plurality of sub-volumes; and

instructions for associating each of the plurality of intermediate nodes with a sub-volume from the plurality of sub-volumes that is larger than the smallest sub-volume.

47. The computer program product of claim 41, wherein

the instructions for selectively casting a plurality of rays from a 2-D image plane include

subdividing the 2-D image plane into a plurality of sub-planes;

for each of the plurality of sub-planes,

casting four rays from the four corners of the sub-plane and estimating a pixel value at each corner;

calculating a maximum pixel value variation within the sub-plane; and

recursively subdividing the sub-plane into multiple child sub-planes of smaller sizes by casting a ray from the center of the sub-plane until the maximum pixel value variation of the sub-plane is below a predefined imaging error threshold.

48. The computer program product of claim 47, wherein the maximum pixel value variation within the sub-plane is defined as the maximum deviation of pixel values at the four corners of the sub-plane from an average pixel value of the sub-plane.

49. The computer program product of claim 47, wherein the predefined imaging error threshold is modulated by an image rendering speed provided by a user, a distance to an edge of an object embedded in the 3-D dataset, and a difference between a pixel value estimated from an adaptive ray casting and a pixel value estimated from a bi-linear interpolation.

50. The computer program product of claim 41, wherein

the instructions for selecting a subset of the plurality of sub-volumes for interacting with the ray include

identifying a largest sub-volume along the ray path and its corresponding maximum and minimum data values;

checking if the opacity transfer function varies monotonically between the maximum and minimum scalar field values;

if the function does not vary monotonically, recursively

identifying a smaller sub-volume along the ray path and its corresponding maximum and minimum data values; and

checking if the opacity transfer function varies monotonically between the maximum and minimum scalar field values of the smaller sub-volume; and

if the function does vary monotonically, calculating the amount of ray energy reflected by the sub-volume during its interaction with the ray.

51. The computer program product of claim 50, wherein two lookup tables are constructed for the opacity transfer function such that a forward lookup table contains the data value difference to a nearest local extreme of the opacity transfer function along the data value increasing direction and a backward lookup table contains the data value difference to a nearest local extreme of the opacity transfer function along the data value decreasing direction.

52. The computer program product of claim 51, wherein if the maximum data value of the sub-volume is smaller than the summation of the minimum data value of the sub-volume and its corresponding data value difference stored in the forward lookup table or the minimum data value of the sub-volume is larger than the difference between the maximum data value of the sub-volume and its corresponding data value difference stored in the backward lookup table, the opacity transfer function varies monotonically between the minimum and maximum data values.

53. The computer program product of claim 41, wherein

the instructions for estimating the ray energy reflected by each sub-volume of the subset include

estimating a maximum energy differential of the sub-volume;

comparing the maximum energy differential against a predefined energy error threshold;

if the maximum energy differential is above the predefined energy error threshold, recursively

selecting a smaller sub-volume along the ray path; and

estimating a new maximum energy differential of the smaller sub-volume; and

if the maximum energy differential is below the predefined energy error threshold, calculating the amount of ray energy reflected by the sub-volume using the illumination model.

54. The computer program product of claim 53, wherein the maximum energy differential depends on the opacity transfer function and the maximum, average, and minimum data values of the sub-volume.

55. The computer program product of claim 53, wherein the amount of ray energy reflected by the sub-volume depends on the length of ray path within the sub-volume, the opacity transfer function within the sub-volume, the average scalar field value of the sub-volume, and the local gradient vector of scalar field within the sub-volume.

56. The computer program product of claim 53, wherein if the sub-volume is a smallest sub-volume comprising 2x2x2 3-D cells, the smaller sub-volume is a 3-D cell within the smallest sub-volume, and if the opacity function does not vary monotonically within the cell, the 3-D cell is further divided into multiple sub-cells until the dimension of a smallest sub-cell reaches the cross-section of the ray.

57. The computer program product of claim 53, wherein if the sub-volume is a smallest sub-volume comprising 2x2x2 3-D cells, the smaller sub-volume is a 3-D cell within the smallest sub-volume, and if the opacity function varies monotonically within the cell, the maximum energy differential of the 3-D cell is estimated by dividing the maximum energy differential of the sub-volume by 2.

58. The computer program product of claim 53, wherein if the sub-volume is a smallest sub-volume comprising 2x2x2 3-D cells, the smaller sub-volume is a 3-D cell within the smallest sub-volume, and if the opacity transfer function varies monotonically within the cell and an iso-surface exists in the 3-D cell, the maximum energy differential of the 3-D cell is calculated using the eight data values at the corners of the 3-D cell and the opacity transfer function.

59. The computer program product of claim 53, wherein the predefined energy error threshold is modulated by an image rendering speed specified by a user, and a zoom factor in the case of parallel projection or a perspective angle and a perspective distance between the image plane and the 3-D dataset in the case of perspective projection.

60. The computer program product of claim 41, wherein

the instructions for estimating pixel values at other locations on the 2-D image plane include

for each location,

selecting four pixel values associated with four ray origins surrounding the location; and

bi-linearly interpolating a pixel value at the location using the four pixel values.

61. A method for generating 2-D images of a 3-D object represented by a sampled 3-D dataset, comprising:

fragmenting the sampled 3-D dataset into a plurality of sub-volumes of different sizes, each sub-volume associated with a set of data value parameters characterizing the data value distribution within the sub-volume;

selectively casting a plurality of rays from a 2-D radiation plane towards the plurality of sub-volumes, each ray having an initial ray energy and a cross-section and launched in a predefined direction with respect to the 2-D radiation plane;

selectively generating a plurality of pixel values at a plurality of locations on a 2-D image plane, each pixel value characterizing an amount of ray energy associated with at least one ray that is reflected by a subset of the plurality of sub-volumes; and

estimating pixel values at other locations on the 2-D image plane using a subset of the plurality of pixel values at the plurality of locations on the 2-D image plane.

62. A method for estimating an amount of ray energy reflected by a 3-D dataset representing a 3-D object when a ray of an initial ray energy interacts with the 3-D dataset along its path, comprising:

fragmenting the 3-D dataset into a plurality of sub-volumes of different sizes, each sub-volume associated with a set of data value parameters including a maximum, an average, and a minimum data value within the sub-volume;

defining an opacity transfer function that is dependent upon data values of the 3-D dataset, the opacity transfer function having a plurality of local extrema at different data values;

constructing a forward lookup table and a backward lookup table, each entry of the forward lookup table containing a data value difference to a local extremum of the opacity transfer function along a data value increasing direction and each entry of the backward lookup table containing a data value difference to a local extremum of the opacity transfer function along a data value decreasing direction;

selecting a set of sub-volumes along the ray path such that the maximum data value of a selected sub-volume is smaller than the summation of the minimum data value of the selected sub-volume and its corresponding data value difference in accordance with the

forward lookup table or the minimum data value is larger than the difference between the maximum data value and its corresponding data value difference in accordance with the backward lookup table;

for each sub-volume of the selected set, estimating an amount of ray energy reflected by the sub-volume in a predefined reflection direction in accordance with a predefined illumination model; and

summing the amount of ray energy reflected by each sub-volume of the selected set together as the amount of ray energy reflected by the 3-D dataset.

63. A method for determining the monotonicity of a function that varies with a variable between a first variable value and a second variable value, comprising:

identifying a first local extremum that is closest to and larger than the first variable value;

comparing the first local extremum with the second variable value;

identifying a second local extremum that is closest to and smaller than the second variable value; and

comparing the second local extremum with the first variable value.

64. The method of claim 63, wherein the function is monotonic between the first variable value and the second variable value if either the first local extremum is larger than the second variable value or the second local extremum is smaller than the first variable value.

65. A data structure for representing an image volume, comprising:

constructing an octree, the octree having a root node, a plurality of intermediate nodes, and a plurality of leaf nodes, and each non-leaf node having eight child nodes;

associating the image volume with the root node of the octree;

fragmenting the image volume into eight sub-volumes and associating a sub-volume with one child of the root node; and

recursively fragmenting a sub-volume into eight smaller sub-volumes and associating a smaller sub-volume with one child of the node associated with the sub-volume.

66. An adaptive direct volume rendering system, comprising:

a plurality of hosts, each host having one or more central processing units for executing programs;

an user interface for receiving a plurality of volume rendering parameters;

a data storage device for storing an image volume; and

a plurality of adaptive direct volume rendering engines, each engine running on at least one of the plurality of hosts.

67. The system of claim 66, wherein the image volume is partitioned into multiple sub-volumes and each engine processes at least one sub-volume and generates at least one sub-image.

68. The system of claim 67, wherein at least one of the plurality of hosts gathers a plurality of the sub-images together and creates an image on a display or output device.